

Modelling Epistemic Systems

André C. R. Martins
GRIFE-EACH, Universidade de São Paulo
São Paulo, SP, Brazil

July 24, 2012

Abstract

In this Chapter, I will explore the use of modeling in order to understand how Science works. I will discuss the modeling of scientific communities, providing a general, non-comprehensive overview of existing models, with a focus on the use of the tools of Agent-Based Modeling and Opinion Dynamics. A special attention will be paid to models inspired by a Bayesian formalism of Opinion Dynamics. The objective of this exploration is to better understand the effect that different conditions might have on the reliability of the opinions of a scientific community. We will see that, by using artificial worlds as exploring grounds, we can prevent some epistemological problems with the definition of truth and obtain insights on the conditions that might cause the quest for more reliable knowledge to fail.

1 Introduction

The classical description of the work of a scientist is one that many of us would like to believe true. According to it, scientists are as close as reasonably possible to believe to selfless individuals who pursue knowledge using always the best means available to them. They propose hypothesis and theories according to what they feel would describe all relevant data better. And, from those models, they draw predictions that are tested under very strict conditions. Finally, still according to our idealized view, it is the agreement between observations and those predictions that dictate which theories are better accepted. Conditions of beauty, like simplicity and symmetry, can be invoked when deciding between equally well suited theories, depending on who is telling the tale. But that is it, and no other considerations should be included.

Obviously, the real life is not so simple and scientists do defend their own ideas because it is their own or because they like it better. And, as in any human activity, one will find individuals whose agenda is based only in self-interest. However, as long as community moves in the right direction and eventually corrects any mistakes, we could, in principle, consider these problems as just perturbations to an otherwise reasonably precise description of the scientific

enterprise. In particular, the amazing advances of knowledge in many areas, where old ideas have been replaced by newer and better ones, seems to lend a lot of credit to that idealized description. At the very least, as a good approximation for the system as a whole, even if it fails in individual cases.

And, indeed, we do have controls to avoid the more serious problems. Replication (or reproducibility) [27, 69, 56] is considered fundamental by most and with good reason. Once an independent researcher obtains the same results, chances of error in the first report are obviously smaller. But, most important of all, if the groups are truly independent and the second group has nothing to gain from either a positive or a negative result, the chances that the first group published fraudulent results basically disappear. This makes the problem of self-serving scientists much less important than it might otherwise be. It also points out to another central feature of the scientific endeavor that is sometimes neglected. That is, the social aspect of Science.

As in many human activities, the outcomes of scientific research are nowadays a social product of a community of scientists. Theory and measurement are often done by different individuals, except in cases where just a fact, a simple idea, or a property of some material or drug are tested. Even in these cases, there is often a group behind the results, as many fields have become so complex that the expertise of many different researchers is often needed. Main examples of this are the "Big Science" projects, such as the mapping of the human genome or the hunt for the Higgs boson. You have different people taking care of different aspects of the problem, some building the equipments, others operating them, a group specializing in collecting the data, more people to store it properly, others to interpret it and so on. Even in smaller projects, like a simple simulation that one can program alone, we often use software developed by a third party, software we trust to do what we are told it does. We depend on the quality of the information and work performed by others in most of the research done nowadays. And, while it might be reasonable, under some circumstances, to assume that most of the information we use were produced in a competent and honest way, it is certain that we can not claim it for all the information[9, 10]. Despite what we are taught in school, arguments of authority are actually unavoidable [28] and we have to understand the effects this might have.

As soon as we realize that we have come to depend on the results of others even to make our own research, we have to realize we might risk running into epistemic trouble. While it is clear that some social aspects are clearly positive (as, obviously, replication), social influence could lead some people away from the best answer. That means that understanding which social aspects can help the reliability of a result and which ones are more likely to be detrimental is an important task. And we must acknowledge that there are different accounts of the scientific process that do not describe it in such beautiful colors, like, per example, the Strong Program in the Sociology of Scientific Knowledge [3, 72, 48, 5]. That program defends that Science is just another human social activity. And, as such, not more reliable than any other descriptions of the world obtained by other methods. While scientific advances seem to contradict

that in a very strong way, one must admit that some of the criticism might be correct and that there might be ways to make the whole scientific activity even more reliable and less prone to error. In a time where no single human can check the correctness of his whole field alone, understanding the consequences of social interaction and errors can be fundamental to the future quality of the Science we will make.

Furthermore, there are fields in which published results have been consistently shown to be wrong[38, 4, 78], exactly due to the way the Science and publications are structured. Biases towards publishing only some types of results exist and are very prejudicial. It is also already known that replication is a far less common practice than it would be expected in areas as different as Marketing [35, 18] and Medicine [37]. These findings make it very clear that it is imperative that we study and understand well the effects of how the way Science is structured and how scientists interact might affect the quality of our conclusions.

Of course, trying to figure how reliable an answer is or if we can really say we know something about the world is an old philosophical problem that we will not answer so simply. However, the question of reliability of results might be particularly well suited to be analyzed with the tools of simulation and agent based models. Simulating a community of scientists will not answer the deepest problems of Skepticism [44]. However, if we assume that it is possible to know something about the world, agent based models of Science can help us explore which strategies and behaviors are more likely to cause problems or to get us closer to the best answers. This happens not only because we can explore unforeseen consequences of interaction rules that we wouldn't see in models built using just human language, but also because we can, as programmers, choose which theory is the best one in the artificial world of our agents. And therefore, in principle, check the effects of different social structures on the choice of the best theory. Of course, any limitations that apply to current models of social interaction will apply here fully and the results must be analyzed with the typical doubt we associate with tentative new ideas. But we can get some advice on where it might be more likely that a change in current practices will lead to better outcomes.

1.1 Existing models for describing scientific practices

In this paper, I defend one specific use for agent models, one that has seen a small number of papers up to now. However, it is interesting to acknowledge that the applications of computational and mathematical modeling to communities of scientists is an area that has been seen a steady increase in the interest and number of papers. Different researchers have dedicated themselves to different problems, most of them tending towards describing the society of scientists. Therefore, it makes sense to provide a very brief and very incomplete review of those applications. Those works can be divided in three major areas: methods to measure the quality and impact of a scientist work; studies of the structure of Science, as per example, the networks of authors and their papers; and the

modeling of knowledge, either just looking at its diffusion, or at the acceptance of new and better knowledge. It is particularly worth noting that recently, there was a full edition of the Journal of Artificial Societies And Social Simulation [15] as well as a book [79], both dedicated to simulation of Science.

Measuring the impact of the work of each of us can be an important task. It has significant influence on our careers and scientific policy makers could use such this information when deciding how to divide the funding [93]. It is no surprise therefore that attempts at providing quantitative measurements for the relevance of a scientist work have become popular and several different proposals now exist. They include the now unavoidable (despite the biases it introduces) H-index [32], initially proposed to avoid the problems of measurement brought by total number of citations, that could be a very large number even if the scientist had just collaborated in one very well cited article. Several other measurements were also proposed [16, 23] and many others are probably being investigated. One interesting aspect of the problem that still needs to be addressed, however, is the impact these indexes have in scientific production. Scientists are usually reasonably intelligent individuals and, therefore, they might change the way they work in order to have better values at the measurement they are evaluated with, instead of worrying only about making better Science.

On the evolution of scientific fields, the ability to measure and quantify production can be both important and interesting in itself [17]. The existence of big databases of published material has allowed a large number of studies on structure that is formed by articles and researchers. There are models describing the networks that are formed when scientists cite others [24, 6, 76, 75, 77] as well as the networks of co-authorship [66, 65, 2], and even of articles connected by the same specific area [31]. This mapping effort allows us to have a much clearer view of how Science is done. Also, the existence of these studies and the data they are based on provides us with indirect measures of social influences among researchers. A complete agent model of scientific activities should either use or explain the structures that were observed; however, we are still far from obtaining such a model, the same way that we are far from having a really reliable and complete model of how people influence each other choices in any field of human behavior. Existing models, as we will see, are still simplistic versions built to address just a few questions and not really general enough.

Still, more recent years have seen the appearance of the first models of idea diffusion in scientific communities, including approaches that make use of the Master equation [45] and population dynamics [92] instead of agent models. Among the agent models, several just try to describe the statistical features observed in the studies that measured number of papers and how they are connected by proposing simple rules of influence [26, 82, 83, 67] such as, per example, copying the citations of read papers. Attempts at introducing more sophisticated agents, with better cognition, also exist [63]. Many general ideas not completely developed have also been proposed as possible basis for future models [12, 70, 71], including general ideas about the practice inspired by the works of philosophers of Science[14].

Of particular interest to the view defended here are models for the diffusion of ideas and opinions that try to answer whether our current practices are good enough or whether we should work to change some of them. One first model used the Bounded Confidence model [11, 29] of Opinion Dynamics [8, 21, 87, 29, 19, 54, 60] as basis, exploring if agents would approach a true value of a parameter given that some of the agents were truth-seekers, who had a tendency to move towards the correct result [30]. However, by adopting a purely continuous version of opinion, that model, while very interesting, is not enough to describe the choice of theories. The effects of peer review [84, 86] were also suggested as an important field to explore by the use of models, since many of the proposals for change in the peer review process have not been explored with the necessary depth. In the same line, discussions of the possible use of modelling to make Science more reliable already exist [94]. However, a strong push forward in developing these models is still to be observed.

Another promising line of investigation on the reliability of Science is the study of the effects of the existence of agents who make claims that are stronger than the actual observations warrant. The main question explored there is if this type of lie could help in the convincing of general public [20, 22]. Finally, I have also proposed a model to explore acceptance of better theories [59], one that was able to show that the idea that Science advances thanks to the retirement of old scientists [46] might be true. In the Section 3, we will discuss it and I will show why I believe it can provide the basis for a more sophisticated and maybe even a little less artificial description of the scientific enterprise. That study was based on an Opinion Dynamics model that distinguishes choices and strength of opinion [54, 55] and is also based on notions of rationality. A model working on the same problem, where actual publications were introduced as the means that scientists use to communicate their results was also recently proposed with some very promising results [85].

2 Epistemology and Modelling

Before presenting a model, it makes sense to discuss a few issues regarding the applicability of mathematical and computational modeling to the epistemic problem of finding better ways to obtain reliable knowledge. The first question worth mentioning, but too complex to be discussed properly here, is exactly the meaning of knowing [80]. Traditional accounts define that one knows something when the person has a justified true belief about that something, but there are known problems with that definition [25].

As mentioned above, one interesting feature of artificial worlds is that, since the programmer is playing the part of a creator god, we can circumvent, for the agents, any of the arguments of Skepticism about the possibility of knowing something about the world. Even softer skeptical arguments, like the problem of induction [36, 33] can be either dismissed by constructing a world where the rules don't change, or, eventually, one could investigate what would happen if the future did not exactly follow the same laws as the past. By controlling

the laws of the artificial universe, different scenarios and their impact on the learning of the agents can be much better investigated than in the real world, where we can only know our theories, but not be always certain about which of them are really better. Of course, this means developing better models for the dynamics of scientific opinions, but the possibility that those models can be quite helpful at exploring epistemic problems should be clear.

In the next Section, I will describe a model that I believe has the potential to be built upon so that we can arrive at a more realistic description of the problem. But, before that, it is good to debate the fundamental ideas behind it and how they relate to current developments in Epistemology. One line of reasoning in Epistemology claims that all we can really know about the world are probabilities that our ideas are correct. When making any new observation, those probabilities must change according to what we have learned, by obeying the Bayes Theorem. This idea can be called Bayesian Epistemology [41, 42, 7, 88]. This is a normative theory, in the sense that it claims that this is the best way to make inferences, and not necessarily the way we actually perform them. The claim that scientists behave and choose the theories they support in a way similar to the specifications of Bayesian methods is known in Philosophy as Confirmation Theory [13, 39, 40, 34, 49, 7]. While it is known that humans are actually bad probabilists [73], there is some evidence that scientists do not deviate so much from a Bayesian point of view [74]. Interestingly, not only scientists, but there is a growing body of evidence that the way people reason about problems and theories is actually close to Bayesian analysis [89, 43], if not exactly so. The so-called error in human probabilistic reasoning can actually be explained by assuming that our brains work with more complex models of reality than what was supposed in the laboratory tests [50, 52].

Therefore, we can use Bayesian models as basis for the behavior of our agents [53], even if only as a reasonable first approximation. Such an approach has already produced a number of Opinion Dynamics models [1, 68, 47, 54, 55, 57, 91, 62, 81, 58, 59, 61, 60] and it was also used to show why a non-specialized reader should only come to trust a scientific model in cases where there might be errors and deception from the part of the authors, when the results are replicated by a third party [51, 69, 56].

3 A Model for the spread of a new theory

Assume each agent must choose between two theories, A and B . This choice is a discrete choice between those alternatives, but agents can have a strength of opinion, represented by the probability p_i that each artificial scientist i assigns to the possibility that A is the best description. If $p_i > 0.5$, agent i believes A is the best option, otherwise, it is B . This model is basically the “Continuous Opinions and Discrete Actions” model (CODA model), previously used to explore the emergence of extremism in artificial societies [54, 55]. CODA is based on the idea that agents believe that, if A is the best choice, each of the neighbors of an agent, located in a given social network [64, 90], will have a probability $\alpha > 0.5$

of choosing A (and similarly, for B). In this context, it is easier to work with a transformation of variables, given by the quantity

$$\nu_i = \ln \frac{p_i}{1 - p_i}.$$

Here, if $\nu_i > 0$, we have $p_i > 0.5$ and, therefore, a subjective belief in favor of A ; if $\nu_i < 0$, the agent chooses B . By applying Bayes Theorem, we obtain a very simple update rule for ν_i , when agent i observes the choice of its neighbor j , given by

$$\nu_i(t+1) = \nu_i(t) + \text{sign}(\nu_j) * a,$$

where a is a step size that depends on how likely the agents believe it is that their neighbors will be correct, that is, it is a function of α . If we renormalize the update rule, by using $\nu_i^* = \nu_i/a$ instead, we will have

$$\nu_i^*(t+1) = \nu_i^*(t) + \text{sign}(\nu_j^*), \quad (1)$$

making it clear that the value of a is irrelevant to the dynamics of choices, since that dynamics depends only on the signs of ν_i (or ν_i^*). The update of the opinions is asynchronous.

In order to introduce the influence of observations in this problem, a proportion τ of the scientists actually perform experiments. This experiments can have a stronger or weaker influence on this agents, measured by the probability ρ and experimenter will, at each interaction, observe the Nature, instead of being influenced by a neighbor. In here, we assume experiments always provide the same answer, agreeing with A and that the agent will update its opinion by the same amount it would from an interaction with a neighbor. Notice that a stronger effect can easily be introduced simply by allowing ρ to be larger and that means no new parameter needs to be introduced here.

It was observed that, unless τ is reasonably large, even when experimentalists are very weakly influenced by social effects (ρ close to 1.0), they fail to convince the whole population and clusters of agents preferring the worst theory survive. This effect was particularly strong when theory A was new and, therefore, started with a small proportion of supporters, indicating that the old view would survive simply from social effects and the strength of opinion of its old supporters. Interestingly, when retirement was introduced in the system, with agents replaced by new ones with moderate initial opinions, it became much easier for new and better ideas to invade, confirming Kuhn's notion that Science would advance and accept new paradigms due to the death of the old scientists [46].

Of course, those results were obtained with a very simplified version of the scientific activity. This was intentional, in the old Physics tradition of starting at the very simplest model one can imagine and later add new features, as they become needed. The idea is to understand which details are responsible for which features of the whole problem. As such, expanding the model from its several approximations is a natural next step. For example, in the model, social

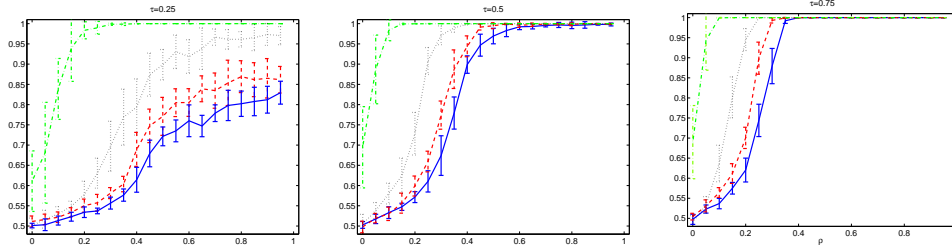


Figure 1: Proportion of agents aligned with the external field (correct scientists) as a function of ρ . Every realization had initial conditions where each agent had 50% chance of supporting either theory. The solid line corresponds to $\lambda = 0.0$, the dashed one to $\lambda = 0.1$, the dotted one to $\lambda = 0.25$, and the dash-dotted one to $\lambda = 0.5$. (*Left panel*) $\tau = 25\%$. (*Middle panel*) $\tau = 50\%$. (*Right panel*) $\tau = 75\%$.

influence happened only between peers and they were all located on a non-realistic square lattice. No publications existed in the model and Nature always gave the same answer, meaning one theory was clearly superior to the other. And, while people would reinforce their ideas given a certain neighborhood, all agents were supposed to be honest, in the sense that they would always choose the theory they really believed more likely to be truth. Despite all this, the model is very easy to expand. By changing the likelihoods agents assign to the choices of others, we could introduce more realistic versions of social influence. Researchers could do papers, adding a new layer to the problem, and the quality of those papers could depend on the correctness of their argument. And, of course, more realistic networks can be trivially introduced.

In order to illustrate these differences, new cases for small-world networks were run specifically for this paper. The results are shown in Figure 1. All cases shown correspond to a initial square lattice with 32^2 agents, where each link was randomly rewired with probability λ . Each point is the result of the average of 10 different realizations of the same problem and the bars correspond to the standard deviation of the observed values. Initial conditions were such that at first both theories had equal amount of support and the system was left to interact for $t = 200$ average interactions per agent. Each panel shows the results for different values of λ for a specific proportion τ of experimenters. We can see that as the small-world effect becomes more important (larger rewiring λ), it becomes easier for the system to reach agreement on the best theory. However, the problem with the continuing existence of groups supporting the worse theory remains for the cases where the importance of experiment is not large enough, confirming, at least, qualitatively, the results previously obtained.

4 Conclusion

We have seen that mathematical and computational modeling of scientific agents can be used as a tool to help us understand under which conditions scientific results are more reliable. While it is quite true that the models so far proposed are very simple, we have seen that they can already capture some of the features of the scientific enterprise, such as the need for solid and reliable experimental data. If experiments are not very convincing, we would be in the situation where ρ is small and, therefore, social influence can indeed become the major force in the system, as suggested in the Strong Program. However, as soon as experiments become more important and we have a better connected world, the case for the better theory becomes strong enough to convince everyone. We should also notice that the same model was able to show the importance of the retirement of scientists, as a way to allow new better ideas to spread more easily.

As such, I would like to encourage the modeling community to investigate models for scientific activity. And by that I mean not only from the sociological (and also interesting) point of view of measuring our activities, but also from the point of view of suggesting better practices. While doubt will remain when a suggestion comes from a model, the tools we have make the exploration of the consequences more reliable than simple human intuition and spoken arguments. Models of Science can and should help us make research an even better tool for understanding the world.

5 Acknowledgements

The author would like to thank Fundação de Amparo à Pesquisa do Estado de São Paulo (FAPESP), under grant 2011/19496-0, for the support to this work.

References

- [1] A. V. Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3):797–817, 1992.
- [2] A.L. Barabási, H. Jeong, Z. Neda, E. Ravasz, A. Schubert, and T. Vicsek. Evolution of the social network of scientific collaborations. *Physica A*, 311:590–614, 2002.
- [3] Barry Barnes and David Bloor. Relativism, rationalism, and the sociology of knowledge. In Martin Hollis and Steven Lukes, editors, *Rationality and Relativism*, pages 21–47. MIT Press, 1982.
- [4] C. Glenn Begley and Lee M. Ellis. Drug development: Raise standards for preclinical cancer research. *Nature*, 483:531533, 2012.
- [5] David Bloor. *Knowledge and Social Imagery*. University of Chicago Press, 1991.

- [6] Katy Börner, Jeegar T. Maru, and Robert L. Goldstone. The simultaneous evolution of author and paper networks. *Proc. Natl. Acad. Sci. U.S.A.*, 101:5266–5273, 2004.
- [7] L. Bovens and S. Hartmann. *Bayesian Epistemology*. Oxford University Press, 2003.
- [8] C. Castellano, S. Fortunato, and V. Loreto. Statistical physics of social dynamics. *Reviews of Modern Physics*, 81:591–646, 2009.
- [9] Erika Check and David Cyranoski. Korean scandal will have global fallout. *Nature*, 438:1056–1057, 2005.
- [10] Mildred K. Cho, Glenn McGee, and David Magnus. Lessons of the stem cell scandal. *Science*, 311:614–615, 2006.
- [11] G. Deffuant, F. Amblard, and T. Weisbuch, G. and Faure. How can extremism prevail? a study based on the relative agreement interaction model. *JASSS-The Journal Of Artificial Societies And Social Simulation*, 5(4):1, 2002.
- [12] Jim Doran. Two outline models of science: Ams and hams. *JASSS-The Journal Of Artificial Societies And Social Simulation*, 14(4):5, 2011.
- [13] J Earman. *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. MIT Press, 1992.
- [14] Bruce Edmonds. A brief survey of some relevant philosophy of science. *JASSS-The Journal Of Artificial Societies And Social Simulation*, 14(4):7, 2011.
- [15] Bruce Edmonds, Nigel Gilbert, Petra Ahrweiler, and Andrea Scharnhorst. Simulating the social processes of science. *JASSS-The Journal Of Artificial Societies And Social Simulation*, 14(4):14, 2011.
- [16] L. Egghe. Theory and practise of the g-index. *Scientometrics*, 69:131–152, 2006.
- [17] L. Egghe and R. Rousseau. *Introduction to Infometrics: quantitative methods in library, documentation and information science*. Elsevier, 1990.
- [18] H. Evanschitzky, C. Baumgarth, R. Hubbard, and J. S. Armstrong. Replication research in marketing revisited: A note on a disturbing trend. *Journal of Business Research*, 60:411–415, 2007.
- [19] S. Galam. Local dynamics vs. social mechanisms: A unifying frame. *Europhysics Letters*, 70(6):705–711, 2005.
- [20] S. Galam. Opinion dynamics, minority spreading and heterogeneous beliefs. In B. K. Chakrabarti, A. Chakraborti, and A. Chatterjee, editors, *Econophysics & Sociophysics: Trends & Perspectives.*, pages 363–387. Wiley, 2006.

- [21] S. Galam, Y. Gefen, and Y. Shapir. Sociophysics: A new approach of sociological collective behavior: Mean-behavior description of a strike. *J. Math. Sociol.*, 9:1–13, 1982.
- [22] Serge Galam. Public debates driven by incomplete scientific data: The cases of evolution theory, global warming and h1n1 pandemic influenza. *Physica A*, 17(1):3619–3631, 2010.
- [23] Serge Galam. Tailor based allocations for multiple authorship: a fractional gh-index. *Scientometrics*, 89(1):365–379, 2011.
- [24] E. Garfield, I. H. Sher, and R. J. Torpie. *The Use of Citation Data in Writing the History of Science*. Institute for Scientific Information, Philadelphia, 1964.
- [25] E. Gettier. Is justified true belief knowledge? *Analysis*, 23:121–123, 1963.
- [26] Nigel Gilbert. A simulation of the structure of academic science. *Sociological Research Online*, 2(2), 1997.
- [27] J. Giles. The trouble with replication. *Nature*, 442:344–347, 2006.
- [28] John Hardwig. Epistemic dependence. *Journal of Philosophy*, 82(7):335–349, 1985.
- [29] R. Hegselmann and U. Krause. Opinion dynamics and bounded confidence models, analysis and simulation. *Journal of Artificial Societies and Social Simulations*, 5(3):3, 2002.
- [30] Rainer Hegselmann and Ulrich Krause. Truth and cognitive division of labour first steps towards a computer aided social epistemology. *Journal of Artificial Societies and Social Simulation*, 9(3):10, 2006. <http://jasss.soc.surrey.ac.uk/9/3/10.html>.
- [31] Mark Herrera, David C. Roberts, and Natali Gulbahce. Mapping the evolution of scientific fields. *PLoS ONE*, 5(5):e10355, 2010.
- [32] J.E. Hirsch. An index to quantify an individuals scientific research output. *PNAS*, 102:16569–16572, 2005.
- [33] C. Howson. *Hume’s Problem: Induction and the Justification of Belief*. Oxford, Oxford University Press, 2003.
- [34] C. Howson and P. Urbach. *Scientific Reasoning: The Bayesian Approach*. Chicago, Open Court, 1993.
- [35] R. Hubbard and J. S. Armstrong. Replications and extensions in marketing: Rarely published but quite contrary. *International Journal of Research in Marketing*, 11:233–248, 1994.

- [36] D. Hume. *An Inquiry Concerning Human Understanding*. Oxford, Clarendon Press, 1777.
- [37] J. P. A. Ioannidis. Contradicted and initially stronger effects in highly cited clinical research. *Journal of the American Medical Association*, 294:218–228, 2005.
- [38] J. P. A. Ioannidis. Why most published research findings are false. *PLoS Med*, 2(8):e124, 2005.
- [39] R. C. Jeffrey. *The Logic of Decision*. Chicago, University of Chicago Press, 1983.
- [40] R. C. Jeffrey. *Subjective Probability: The Real Thing*. Cambridge, Cambridge University Press, 2004.
- [41] M. Kaplan. *Decision Theory as Philosophy*. Cambridge University Press, 1996.
- [42] M. Kaplan. Decision theory and epistemology. In P. K. Moser, editor, *The Oxford Handbook of Epistemology*. Oxford University Press, 2002.
- [43] C. Kemp, J. B. Tenenbaum, S. Niyogi, and T. L. Griffiths. A probabilistic model of theory formation. *Cognition*, 114(2):165–196, 2010.
- [44] P. Klein. Skepticism. In P. K. Moser, editor, *The Oxford Handbook of Epistemology*. Oxford University Press, 2002.
- [45] Pawel Kondratiuk, Grzegorz Siudem, and Janusz A. Holyst. Analytical approach to model of scientific revolution. arxiv preprint 1106.0438, 2011.
- [46] T. Kuhn. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press, 1962.
- [47] D. Lane. Is what is good for each best for all? learning from others in the information contagion model. In W.B. Arthur, S.N. Durlauf, and D. Lane, editors, *The Economy as an Evolving Complex System I*, pages 105–127. Santa Fe: Perseus Books, 1997.
- [48] Bruno Latour and Steven Woolgar. *Laboratory Life: The Construction of Scientific Facts*. Princeton University Press, 1986.
- [49] P. Maher. *Betting on Theories*. Cambridge, Cambridge University Press, 1993.
- [50] André C. R. Martins. Adaptive probability theory: Human biases as an adaptation. Cogprint preprint at <http://cogprints.org/4377/>, 2005.
- [51] André C. R. Martins. Deception and convergence of opinions. *JASSS-The Journal Of Artificial Societies And Social Simulation*, 8(2):3, 2005.

- [52] André C. R. Martins. Probabilistic biases as bayesian inference. *Judgment And Decision Making*, 1(2):108–117, 2006.
- [53] André C. R. Martins. A bayesian framework for opinion updates. arXiv:0811.0113v1, 2008.
- [54] André C. R. Martins. Continuous opinions and discrete actions in opinion dynamics problems. *Int. J. of Mod. Phys. C*, 19(4):617–624, 2008.
- [55] André C. R. Martins. Mobility and social network effects on extremist opinions. *Phys. Rev. E*, 78:036104, 2008.
- [56] André C. R. Martins. Replication in the deception and convergence of opinions problem. *JASSS-The Journal Of Artificial Societies And Social Simulation*, 11(4):8, 2008.
- [57] André C. R. Martins. Bayesian updating rules in continuous opinion dynamics models. *Journal of Statistical Mechanics: Theory and Experiment*, 2009(02):P02017, 2009. arXiv:0807.4972v1.
- [58] André C. R. Martins. A middle option for choices in the continuous opinions and discrete actions model. *Advances and Applications in Statistical Sciences*, 2:333–346, 2010.
- [59] André C. R. Martins. Modeling scientific agents for a better science. *Adv. Compl. Sys.*, 13:519–533, 2010.
- [60] André C. R. Martins. *Social Physics Catena (No.3)*, chapter Discrete Opinion models as a limit case of the CODA model, pages 146–157. Science Press, 2012. arXiv:1201.4565v1.
- [61] André C. R. Martins and Cleber D. Kuba. The importance of disagreeing: Contrarians and extremism in the coda model. *Adv. Compl. Sys.*, 13:621–634, 2010. arXiv:0901.2737.
- [62] André C. R. Martins, Carlos de B. Pereira, and R. Vicente. An opinion dynamics model for the diffusion of innovations. *Physica A*, 388:3225–3232, 2009.
- [63] Isaac Naveh and Ron Sun. A cognitively based simulation of academic science. *Computational and Mathematical Organization Theory*, 12:313–337, 2006.
- [64] M. Newman, A.-L. Barabasi, and D. J. Watts. *The Structure and Dynamics of Networks*. Princeton, Princeton University Press, 2006.
- [65] M. E. J. Newman. Scientific collaboration networks: I. network construction and fundamental results,. *Phys. Rev. E*, 64:016131, 2001.
- [66] M. E. J. Newman. The structure of scientific collaboration networks. *Proc. Natl. Acad. Sci. U.S.A.*, 98:404–409, 2001.

- [67] M. E. J. Newman. The first-mover advantage in scientific publication. *Europhys. Lett.*, 86:68001, 2009. arxiv:0809.0522v1.
- [68] André Orléan. Bayesian interactions and collective dynamics of opinion: Herd behavior and mimetic contagion. *Journal of Economic Behavior and Organization*, 28:257–274, 1995.
- [69] V. Palmer. Deception and convergence of opinions part 2: the effects of reproducibility. *Journal of Artificial Societies and Social Simulation*, 9(1):14, 2006.
- [70] Sergey Parinov and Cameron Neylon. Science as a social system and virtual research environment. *JASSS-The Journal Of Artificial Societies And Social Simulation*, 14(4):10, 2011.
- [71] Nicolas Payette. For an integrated approach to agent-based modeling of science. *JASSS-The Journal Of Artificial Societies And Social Simulation*, 14(4):9, 2011.
- [72] Andrew Pickering. *Constructing Quarks: A Sociological History of Particle Physics*. Edinburgh University Press, 1984.
- [73] S. Plous. *The Psychology of Judgment and Decision Making*. New York, McGraw-Hill, 1993.
- [74] S. James Press and Judith M. Tanur. *The Subjectivity of Scientists and the Bayesian Approach*. Wiley, 2001.
- [75] Walter Quattrociocchi and Frederic Amblard. Emergence through selection: The evolution of a scientific challenge. arXiv preprint: 1102.0257, 2011.
- [76] Filippo Radicchi, Santo Fortunato, Benjamin Markines, and Alessandro Vespignani. Diffusion of scientific credits and the ranking of scientists. *Phys. Rev. E*, 80:056103, 2009.
- [77] Fu-Xin Ren, Xue-Qi Cheng, and Hua-Wei Shen. Modeling the clustering in citation networks. *Physica A*, 391:3533–3539, 2012. arXiv preprint 1104.4209v2.
- [78] Daniel Sarewitz. Beware the creeping cracks of bias. *Nature*, 485:149, 2012.
- [79] Andrea Scharnhorst, Katy Börner, and Peter van den Besselaar, editors. *Models of Science Dynamics: Encounters Between Complexity Theory and Information Sciences*. Understanding Complex Systems. Springer, 2012.
- [80] R. K. Shope. Conditions and analyses of knowing. In P. K. Moser, editor, *The Oxford Handbook of Epistemology*. Oxford University Press, 2002.

- [81] Xia-Meng Si, Yun Liua, Fei Xionga, Yan-Chao Zhanga, Fei Ding, and Hui Cheng. Effects of selective attention on continuous opinions and discrete decisions. *Physica A*, 389(18):3711–3719, 2010.
- [82] M.V. Simkin and V.P. Roychowdhury. Copied citations create renowned papers? *Annals Improb. Res.*, 11(1):24–27, 2005.
- [83] M.V. Simkin and V.P. Roychowdhury. A mathematical theory of citing. *Journal of the American Society for Information Science and Technology*, 58(11):1661–1673, 2007.
- [84] Pawel Sobkowicz. Peer-review in the internet age. arXiv:0810.0486v1, 2010.
- [85] Pawel Sobkowicz. Simulations of opinion changes in scientific communities. *Scientometrics*, 87(2):233–250, 2011.
- [86] Flaminio Squazzoni and Károly Takács. Social simulation that ‘peers into peer review’. *JASSS-The Journal Of Artificial Societies And Social Simulation*, 14(4):3, 2011.
- [87] K. Sznajd-Weron and J. Sznajd. Opinion evolution in a closed community. *Int. J. Mod. Phys. C*, 11:1157, 2000.
- [88] William Talbott. Bayesian epistemology. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. The Metaphysics Research Lab, 2008.
- [89] J. B. Tenenbaum, C. Kemp, and P. Shafto. Theory-based bayesian models of inductive reasoning. In A. Feeney and E. Heit, editors, *Inductive reasoning*. Cambridge University Press., 2007.
- [90] F. Vega-Redondo. *Complex Social Networks*. Cambridge University Press, 2007.
- [91] R. Vicente, André C. R. Martins, and N. Caticha. Opinion dynamics of learning agents: Does seeking consensus lead to disagreement? *Journal of Statistical Mechanics: Theory and Experiment*, 2009:P03015, 2009. arXiv:0811.2099.
- [92] NK Vitanov and MR Ausloos. Knowledge epidemics and population dynamics models for describing idea diffusion. In A Scharnhorst, K Börner, and P van den Besselaar, editors, *Models of science dynamics*. Springer, 2012.
- [93] Levent Yilmaz. Toward multi-level, multi-theoretical model portfolios for scientific enterprise workforce dynamics. *JASSS-The Journal Of Artificial Societies And Social Simulation*, 14(4):2, 2011.
- [94] Kevin Zollman. Computer simulation and emergent reliability in science. *JASSS-The Journal Of Artificial Societies And Social Simulation*, 14(15):4, 2011.